

The Real-Time Display of Interferometry Data for Goldstone Radar Astronomy Data Acquisition

S. D. Howard

Communications Systems Research Section

A method for visualizing radar interferometry data in real time has been developed for the Goldstone radar astronomy ranging data-acquisition system. The presentation is similar in appearance to a vector field display or data-based grid. This form was selected to facilitate the recognition of characteristic patterns of local variation in the phase and magnitude of complex elements in a two-dimensional data array. The design emphasized efficiency under the demands of real-time processing and remote monitoring. The interferometry "phase-magnitude" presentation, as it has come to be called, has been used to monitor radar interferometry experiments on three targets, beginning with the asteroid 4179 Toutatis, and continuing with Mars and Mercury.

I. Introduction

The Goldstone radar astronomy ranging system supports ground-based radar imaging. Radar echoes from the surface of a solar system target are mapped by range delay and Doppler frequency. However, delay-Doppler maps of ideal spherical objects are symmetrical about a central axis: the north and south elements are not resolvable without additional information. Radar interferometry, which combines data from more than one receiving station, resolves the north-south ambiguity by adding phase information to the delay and Doppler information. Interferometry products include not only images of surface reflectivity, but also root-mean-square surface roughness and polariza-

tion ratios. With a satisfactory signal-to-noise ratio (SNR) and sufficient data, altitude can be estimated as well.

Before November 1990, the interferometry data used for producing unambiguous radar images were derived during data reduction from recorded complex voltage data. In 1990, signal processing software devised for producing a new interferometry data type in real time was added to the data-acquisition system. The Appendix describes the interferometry or "crosspower" data type. Data of this type could be integrated without losing information, thus reducing data transfer and recording rates in the data-acquisition system and making radar imaging derived

from interferometry possible for the first time on high-bandwidth targets like Mars.

The 1990 Mars observations, the first using real-time interferometry signal processing, relied on postacquisition display to confirm that useful interferometry data had been acquired. The real-time display in use at the time could display only limited numbers of delay-Doppler power spectra; there were no means of displaying interferometry data. Even though the first Mars experiments collected useful data, there were significant advantages to being able to monitor data acquisition in real time. For example, if the range delays of echoes returning to each station in the interferometer had not been correctly aligned, the interferometry data would have been unusable. Early indications of incorrect station delay offsets or system malfunctions would have allowed station personnel to take corrective action during the observation.

The redesign of the software that manages the data-acquisition display created an opportunity to address the problem of displaying interferometry data in real time. The new software architecture was designed to be extensible. One of the purposes for the design of an extensible architecture was to provide a basis for experimentation with scientific data visualization. The development of the interferometry presentation was an experiment with data visualization and a demonstration that new display methods could be added to the real-time display software as planned. Research and development (R&D) is often associated with the beginning of the product life cycle. Since the ranging system is a major science instrument late in the life cycle, this experience also became an opportunity to experiment with the concept of R&D late in the product life cycle.

Several frequently occurring terms are often used in a specialized sense in this article. A "method" is an operation or transformation applied to a data type or object. This usage is intended to be close to the object-oriented usage. A "presentation" is an abstraction that describes how information will be transformed into a visual representation. A presentation may be rendered by hand or by computer. The term "display" is also often used to mean the real-time process in the data-acquisition computing system, implemented in software, that manages the display of incoming data.

II. Design Objectives and Constraints

There were several overall objectives for the design of the interferometry presentation. Minimizing the impact

of the display on real-time performance was important, and so was retaining the capability for remote monitoring. Economical development and implementation were also primary objectives. Because the display would be introduced into a system in the maintenance phase of the life cycle, it was an objective to adapt the design to prevailing usage, with the intention of minimizing user familiarization time and errors. Both the display software and its platform were to be designed to be maintainable and reliable. Since it was also an objective to make the display available as much as possible, the implementing platform was constrained to low-cost, widely available hardware and software so that the physical display could be quickly replaced, if necessary.

A summary of design objectives related to the performance and function of the system is given in Table 1; design objectives related to the use of the system are listed in Table 2; design objectives derived from an analysis of the project situation are given in Table 3. These objectives shaped the synthesis of a final design. Candidate designs that were variations on presentations used at the time by either data reduction or data acquisition were set aside because they fell outside performance or usability objectives.

III. Design Description

A. Overview

The new interferometry presentation came to be called the phase-magnitude presentation for its emphasis on those dimensions of the data. An example is shown in Fig. 1. The graphic is an adaptation of a data-based grid [1]. Interferometry data are structured as a two-dimensional array of complex elements. The major dimensions are range (or delay) and Doppler frequency. Because each complex element also has two dimensions, the graphical problem became one of presenting four data dimensions on a two-dimensional area. The data graphic solved this problem by creating a two-dimensional array or grid of two-dimensional areas. A square area represents the complex plane with the origin at the center for each complex element. Complex elements are arranged in rows and columns for the two major dimensions of the data. Range is on the vertical axis, and Doppler frequency is on the horizontal axis. The structure of the data presentation is similar in form to the structure of the data.

The arrangement of vectors, number of vectors, size, and scaling was also defined by the design. To avoid the visual confusion caused by crossing vectors, the data graphic was designed so that vector areas would not overlap. During the design process, it was realized that only a sub-

set of the data needed to be shown. The design balances the number of vectors that can be displayed on the screen against the size of the individual vectors. Because only a subset of the data can be seen in the data graphic area, the user has the option of selecting the area of interest. The maximum component of any vector appearing in the display area is used for scaling. Since scaling is affected by the maximum complex component value, the selection of the data area also affects scaling.

B. Details

The text presentation of radar configuration information in the first and last lines is shared with all other presentations in the new real-time display software. These standard text areas, showing time, target, and configuration, form a common visual framework for all presentations in the display set. Only the fundamental parameters are shown; other parameters can be derived from these to establish the configuration. To save screen space and display time, only the lower and left axes are shown. The ratio of the sides of the data area is 32/20 units, close to the ratio of the Golden Rectangle. Visual balance and aesthetics were considered in the definition of the presentation.

Color was used to emphasize some graphical elements and deemphasize others. For instance, transmit and receive range code offsets, critical to the management of the real-time system, are shown in red on the display screen for emphasis whenever they are active. They are also clustered at the top center of the presentation for the same reason. Axis labels and tick marks are displayed in cyan, a light graph-paper blue, to create a mild illusion of space between the data graphic and the closely surrounding text.

C. Interpretation

The data-based grid was employed in this design because it emphasizes patterns of local variation in data. In interferometry data, there are two significant features. The most noticeable is an abrupt change in magnitude that marks the return of the radar echo. The most important from the standpoint of interferometry is the appearance of phase fringes, characteristic local patterns of phase change. For the presentation to be useful, viewers should be able to recognize patterns in the data when significant relationships in phase or magnitude exist. When no important relationships between data elements are present, phase and magnitude dimensions of data elements in the graphic should appear to be unrelated.

IV. Early Results

The evaluation of the design took place in three stages. In the first stage, receiver noise was used to evaluate viewer response; an illustration of this is given in Fig. 2. Viewers saw noise as random. Performance, function, and usability were found to be well within design goals. Additional system testing with loop-back data simulating point-source signals would not have been adequate for evaluating the complex phase relationships that were likely to be received from a planetary target. The new display software architecture had anticipated this situation. It had been designed to accommodate the playback of recorded data. Since no interferometry observations were scheduled for several months, playback became the best option for gaining familiarity with the handling of the new display. The second stage began when existing data sets were reviewed to learn what the data would look like when presented in this way, how much data integration would be necessary to see and interpret the results, and whether the policy of coupling the selection of the data area to scaling would be as acceptable as predicted by the design. The third stage, evaluation during an observation, assessed the impact of the display on real-time operations and instrument management. A summary of results from the first real-time observations is given in Table 4.

A. Playback

1. **Mars 1990.** Playback of the Mars 1990 data demonstrated that the phase patterns of usable interferometry data could be recognized in the new presentation. Figure 1 is an example. In the original review of the data with the data reduction display, it was not possible to see whether phase angles in phase bands changed gradually or abruptly, because the data-reduction display was a color display that associated phase angles in 30-deg intervals with one assigned hue. In the phase-magnitude presentation, since phase angles are shown directly, the gradual change in phase around the front of the planet is apparent.

It is an advantage for the viewer to be able to see phase fringes in the display with integration times less than the maximum. The review of playback data established that fringes could be seen with the new presentation at least as soon as they could be seen with the nonreal-time data-reduction display. Figure 1 was produced from data integrated for 5 sec. To give a good indication of phase fringes, the data-reduction review of the same data had been set to a 20-sec integration. This demonstrated that the integration required for effective data representation by the phase-magnitude presentation was within functional goals.

After playing back the first part of the Mars data set, there was some concern about whether it would be possible to see phase patterns farther back than the first two or three range gates on the planet, given the scaling policy initially implemented for this presentation. For example, in Fig. 1, the phasors representing early returns are near maximum length at center frequency gate 64 and in range gates 16 and 17; in range gate 21, the vectors have become points. This concern was resolved after it was noted that an unusually strong initial echo was present in the data. The display was an accurate representation of the given data.

2. Mercury 1992. Several attempts to acquire Mercury interferometry data were made in 1992. On one occasion, data containing echoes were recorded from both stations in the interferometer, but because of a low SNR from the second station, DSS 13, and uncertainty about whether or not both stations were correctly aligned in range, the value of the interferometry portion of the data was doubtful. Playback of the data with the new presentation did not show phase patterns. This is illustrated in Fig. 3. The area of increased magnitude is the returned echo. Phase appeared to become even more random when data integration time was increased. Increasing integration also caused an apparent decrease in the overall change in magnitude between noise and returned signal. Both of these consequences of increased integration indicated that usable interferometry data had not been recorded.

B. Real-Time

1. 4179 Toutatis. The first real-time experience with the ranging interferometry display took place during the observations of asteroid 4179 Toutatis in December 1992. Station offsets for correct range alignment were unknown when the Toutatis interferometry experiment began. Figure 4(a) is the delay-Doppler display showing the initial range alignment. As with the interferometry presentation, range is the vertical dimension and frequency is the horizontal. After the correct station offsets were included in the system, the delay-Doppler display in Fig. 4(b) shows the final range alignment in range gate 41 for both stations. Phase fringes were recognized as soon as the first phase-magnitude display was received; that display is shown in Fig. 5. Interferometry data points from the most distant lobe of Toutatis appear on the left side, between frequency gates 65 and 76.

The option to select regions of the data to be displayed was utilized during this observation. In the 0.5- μ sec range resolution configuration used for Toutatis, a large system artifact was present at dc (frequency gate 64) in all range

gates. When the artifact was shown in the data graphic, surrounding data values were scaled to points. Moving the data display area to the left or right side of dc allowed normal scaling to the maximum data value.

Some experience with discovering and correcting problems during an observation was also obtained when science and engineering personnel at Goldstone and JPL noticed a degradation in the quality of the displayed interferometry data after several minutes of acquisition. Operations personnel began a search for the problem and noticed that a polarization switch had been set incorrectly. The problem was corrected and, subsequently, recorded data were of acceptable quality.

2. Mars 1992–1993. The first interferometry experiments for the 1992–1993 Mars opposition began 3 days after the Toutatis experiment. The range alignment procedure was repeated for Mars because the range alignment offsets for the Toutatis configuration with submicrosecond range resolution were not correct for the Mars configuration, which had a 6.0- μ sec range resolution. The investigator reported that phase fringes were seen in real time during the first experiment, confirming correct range alignment. Figure 6 is an example of data taken later in the Mars series.

In subsequent Mars experiments, episodes of low SNR caused more switching between the delay-Doppler power spectra display (the two-channel Toutatis delay-Doppler displays shown in Fig. 4 are examples of this type) and the new interferometry display than had been expected. Because the delay-Doppler display emphasizes magnitude only, it is the presentation of choice when the SNRs of both stations in the interferometer are low.

3. Mercury 1993. The Mercury experiment used the same range alignment offsets as the previous Mars experiments, so the range alignment procedure was not repeated or checked. Science and engineering staff once again exercised the remote monitoring capability by remotely monitoring the returns from Goldstone from an office at JPL. A low SNR from the second station in the interferometer required long integration times to see the target, longer than the round-trip light time. Because of the low SNR, the acquisition of usable interferometry had been doubtful. However, phase fringes were recognized immediately when the first display was received. Figure 7 shows the first display.

During the observation, a drop in SNR was noticed in the interferometry display. The loss of SNR was pursued at the station, but the situation could not be corrected before

the end of the observation. Figure 8 shows an interferometry display after the drop in SNR. Maximum magnitude, shown on the right side, above the data display area, is a factor of three lower than the maximum shown in Fig. 7; the apparent magnitude difference between noise and echo has decreased in the data graphic, and indications of phase patterns have disappeared.

The next series of observations, during a favorable inferior conjunction in July 1993, showed that phase bands could be seen at least as far back as 17 range gates, in contrast to only 5 range gates in the Mars data shown in Fig. 1. Additional experience will be necessary to see what effects features on the surface of Mercury have on phase patterns.

C. Phase Features

The ability to read and understand detail in the displays is still developing. One experimenter has created the terms "block" and "swirl" for phase features that can be seen in the phase-magnitude interferometry displays. For planetary targets where the target is large when compared to the fringe spacing, blocks and swirls correspond principally to variations in surface altitude. For small bodies, where the target is small relative to the fringe spacing, blocks, or large areas of equal phase, are the expected form. The Toutatis data in Fig. 5, between range gates 43 and 49 and frequency gates 65 and 76, can be thought of as a block. In the 1990 Mars data, an area that has been referred to as a swirl can be followed in Fig. 9 as it moves from right to left on the display. The swirl is first seen in Fig. 9(a), in the area between range gates 10 and 13 and frequency gates 72 and 80. In Fig. 9(b), it has moved to frequency gates 69 through 76, between range gates 9 and 12. In Fig. 9(c), the swirl has moved toward the front of the planet in the area of frequency gates 66 through 74.

V. Conclusions

The early results from the Toutatis, Mars, and Mercury experiments with the phase-magnitude presentation of interferometry data have demonstrated that the design is efficient for real-time use with remote monitoring and that it is effective for presenting interferometry data to viewers. The implicit objective of this development was to improve the management of the radar during observations, making the acquisition of radar interferometry data sets routine rather than sporadic. Experience with the first three targets, summarized in Table 4, indicates that this objective was met. Use of the phase-magnitude display has given radar scientists a different and precise view of phase features. Although the phase-magnitude presentation was

not designed for this purpose, it is being considered for data-reduction work as well.

The design accommodated the objectives summarized in Tables 1 through 3. The emphasis was on economy in performance and implementation. The displays have been frequently used for remote monitoring at JPL over phone lines transmitting data at 9600 bps; they complete in less than 8 sec. The implementation utilized low-cost, widely available microcomputers running widely available terminal emulator packages. The software that supported simple graphics commands for the radar control computer was already available on the system. Design and development time took less than 3 weeks for one software engineer working part time. The software was written to be maintainable and reliable; however, it has not yet required maintenance or modification. The only requested change has been minor, a modification to the way in which the "max" parameter is computed for display annotation.

Because this experimental design is part of an adaptive maintenance activity, the design also incorporated patterns of current use. User selections were simple and similar in structure to selection parameters for other presentation types. The controls were planned to minimize familiarization time and user errors. Experience with all user groups suggested that the display was straightforward and easy to use. The products of the new presentation were also designed to be used in established ways. For example, the circumstances that motivated the goal described in Table 2, that the display translate well to black and white for hard copy, occurred shortly after the first real-time observation with the new display method, when printouts were requested and sent by facsimile to Arecibo Observatory in Puerto Rico. Circumstances again tested the usability of the presentation in black and white during an interval when the station display device was only capable of black and white output.

This design experiment also demonstrated that new display methods could be included in the real-time software while minimizing degradation to the display software structure. The new software architecture was intended to enable engineering experimentation as part of the research and development charter of radar astronomy while maintaining the reliability of a major science instrument. The design of the phase-magnitude presentation for interferometry data reflects both the experimentation and maintenance interests of the project, and suggests that although R&D work is often associated with the early phases of a project, experimentation can also bring new results in the maintenance phase of the life cycle.

Acknowledgments

The author thanks Paul Willis, Carl Franck, and Ray Jurgens. Paul Willis's comments, suggestions, and perseverance improved this article. Carl Franck and Ray Jurgens assisted during system testing. Ray Jurgens has also provided insights and helpful observations while working with and interpreting the new interferometry presentation.

Reference

- [1] E. R. Tufte, *The Visual Display of Quantitative Information*, Cheshire, Connecticut: Graphics Press, 1983.

Table 1. Summary of design objectives derived from the system.

Analysis	Derived design objectives
Real-time system with finite capacity for additional processing.	Minimize computing and transmission overhead of graphics commands.
Remote system testing and monitoring of observations reduces project costs.	Retain remote monitoring capability.
Time to draw standard delay-Doppler display can exceed 30 sec.	Design new display to complete in less than 30 sec on a terminal emulator receiving data from the host at 9600 bps.
Data sent to display every n sec, where n varies with the configuration from subsecond intervals to minutes (an architectural constraint, costly to change).	Stay within the system architecture. Plan to accept data every n sec.

Table 2. Summary of design objectives derived from the users.

Analysis	Derived design objectives
Users frequently make black and white hard copy from printers, faxes, and photocopiers.	Design the data graphics in the presentation to map well to black and white, without losing information.
Conference papers and publications that publish early results often restrict the use of color graphics.	Same as above.
In the past, only monochrome displays have been used in the real-time system.	Begin transition from monochrome screen graphics to color. Use color for emphasis and focusing viewer attention, rather than as a means of encoding data.
Equipment available at Goldstone is more likely to support monochrome displays; suitable equipment to support color is more difficult to acquire.	Be able to fall back to monochrome with minimal software adaptation if equipment availability becomes a problem.
Color vision defect can cause misinterpretation if red and green elements are adjacent.	Avoid designs that could generate red and green adjacent elements.
Graphics showing raw data have been published in the past.	Design publication-quality graphics with attention to good aesthetics.
Past displays omitted time, target, and basic parameters from which the configuration could be reconstructed.	Design within the constraints of the revised display process; include this supporting information in the display.
Observations are episodic; operators must often refamiliarize before each series.	Minimize the number of parameters that must be learned by operators to use the display.
Changes in the system can cause an increase in operator errors during initial uses.	Minimize the amount of adaptation that the operators must make to use the new display. Keep the system familiar.

Table 3. Summary of design objectives derived from the task.

Analysis	Derived design objectives
Front-end digital signal processing hardware undergoing modification.	Limit the scope of the presentation design and implementation.
Software development computer and file system to be removed or relocated.	Same as above.
One software engineer available part-time for development.	Same as above. Plan for development time of less than 1 workmonth.
Good software development practice has improved the reliability and maintainability of the software.	Use good software practice. Design for robustness, defect prevention, and maintainability.

Table 4. Interferometry results for the first three interferometry targets.

Radar target	Date, Universal Time	Significance	Phase fringes observed	Summary
Asteroid 4179 Toutatis	December 13, 1992	First opportunity to use new interferometry display. Opportunity for first Goldstone real-time ranging interferometry display.	Yes	Fringes seen immediately on receipt of initial display. (See Fig. 5.) Used delay-Doppler display for visual range alignment to determine correct range offsets. Positioned display off center to avoid dc system artifact.
Mars	December 15, 1992	First opportunity to use new interferometry display on Mars. Second interferometry target attempted with new display.	Yes	Reports of phase fringes from investigator.
Mercury	March 20, 1993	First opportunity to use new interferometry display on Mercury. Third interferometry target attempted with new display.	Yes	Fringes seen immediately on receipt of initial display. (See Fig. 7.) SNR dropped later; suspected station problem. No opportunity to look for influence of topography on phase patterns because of low SNR.





Fig. 1. The 1990 Mars Interferometry playback.

(a)



(b)




Fig. 4. Delay-Doppler displays of 4179 Toutatis (a) before range alignment and (b) after range alignment.




Fig. 5. The first real-time interferometry display: asteroid 4179 Toutatis.




Fig. 6. Real-time interferometry presentation of Mars data (1992-1993).




Fig. 7. The first real-time interferometry display of Mercury.





Fig. 8. Three-fold decrease in maximum power during Mercury data acquisition.

(a)



(b)




Fig. 9. A "swirl" seen in playback of the Mars 1990 data (a) on the edge of the echo, right-hand side of the display, (b) moving toward the front as the planet rotates, and (c) near the front of the planet.

(c)




Fig. 9 (cont'd)

Appendix

Interferometry Data Type Description

Let $(t_1, \dots, t_n)_{r,c}$ represent a vector of complex voltage samples, t_i taken on a channel $c, c = 0, \dots, C - 1$, where C is the number of complex channels for range gate $r; r = 0, \dots, R - 1$, where R is the number of range gates. The variable n is a power of 2, limited by implementation to be between 64 and 2048.

The expression $(f_1, \dots, f_n)_{r,c}$ represents the result of applying a discrete Fourier transform to $(t_1, \dots, t_n)_{r,c}$.

The complex spectra from selected channels l and m are combined by applying a complex conjugate multiply to corresponding elements from both spectra:

$$X_{r,1 \times m} = (f_{1,r,m}^* f_{1,r,l}, \dots, f_{n,r,m}^* f_{n,r,l})$$

Crosspower spectra for channels l and m are then computed for all range gates r in the given configuration to form the interferometry data structure.

A New Presentation of Complex Voltage Data for Goldstone Radar Astronomy

S. D. Howard

Communications Systems Research Section

A presentation method similar to a vector field display or a data-based grid has been used to visualize complex voltage test data taken over multiple subchannels. Unlike simple plots of all data points on the complex plane, the position of the data in the time series is part of the presentation, providing additional information to aid in fault isolation during readiness testing. The "phase-magnitude" presentation, as it has come to be called, was designed for the demands of real-time data acquisition and processing with remote monitoring over low bandwidth communication links.

I. Introduction

A new real-time presentation of complex voltage data taken over multiple subchannels has been added as a display option to the Goldstone Solar System Radar ranging system. The new voltage data presentation was adapted from a design devised for complex interferometry data [1], referred to as a "phase-magnitude" presentation. It is similar in appearance to a vector field display or a data-based grid [2]. The voltage phase-magnitude presentation was intended to improve the effectiveness of readiness testing by enabling engineers to discover and correct defects before the data from an observation could be degraded or lost during acquisition due to component failures in the special-purpose signal processing hardware.

The ranging data-acquisition system was initially developed as a rapid prototype. Over the past 5 years, the software has undergone several significant upgrades to enhance reliability, add functionality, improve performance, and correct defects. An early example of one of these reengineering episodes is described in [3]. The most recent

upgrade, a replacement of the software that manages the real-time data display, made complex voltage data available to the display and provided an extensible software architecture for experimenting with new methods of visualizing radar data. With the exception of the earliest version of the ranging system (1985) where power profiles were summed from complex voltage data and displayed, prior to the upgrade conjectures about the state of the system during readiness testing were made from integrated power spectra shown in the delay-Doppler display.

Readiness testing in the radar data-acquisition system typically takes the form of local loop-back testing. During a loop-back test, a point-source test signal is injected into the data-acquisition system before the analog-to-digital converters. Incoming test data are then digitized, autocorrelated, coherently summed, transformed to the frequency domain, detected by taking the magnitude squared, incoherently summed, and displayed in real time. The frequency and time delay of the test signal can be set by the operator. Engineers examine the real-time display as